arXiv:1610.01929v1 [cs.GT] 5 Oct 2016

Trial-Offer Markets with Continuation

Pascal Van Hentenryck* Alvaro Flores† Gerardo Berbeglia‡

October 7, 2016

Abstract

Trial-offer markets, where customers can sample a product before deciding whether to buy it, are ubiquitous in the online experience. Their static and dynamic properties are often studied by assuming that consumers follow a multinomial logit model and try exactly one product. In this paper, we study how to generalize existing results to a more realistic setting where consumers can try multiple products. We show that a multinomial logit model with continuation can be reduced to a standard multinomial logit model with different appeal and product qualities. We examine the consequences of this reduction on the performance and predictability of the market, the role of social influence, and the ranking policies.

1 Introduction

With the ubiquity of online market places such as `Amazon` and `iTunes`, there has been increasing interests in understanding and modeling the behavior of such trial-offer markets, where customers sample a product before deciding whether to buy it. These online markets are particularly interesting because of their greater opportunities in shaping the customer experience and their flexibility in exploiting visibility bias Lerman and Hogg (2014); Buscher, Cutrell, and Morris (2009); Maillé et al. (2012); Joachims et al. (2005) and social signals Engstrom and Forsell (2014); Viglia, Furlan, and Ladrón-de Guevara (2014).

Traditionally, such markets have been studied using extensions of Multinomial Logit Models Krumme et al. (2012); Lerman and Hogg (2014); Abeliuk et al. (2015); Van Hentenryck et al. (2015); Daly and Zachary (1978); Talluri and Van Ryzin (2004); Rusmevichientong, Shmoys, and Topaloglu (2010); Rusmevichientong, Shen, and Shmoys (2010), where participants can only sample one product (e.g., listening to a song) before deciding whether to buy the product. Multinomial Logit Models have been studied in marketing for several decades Luce (1965), but the addition of social signals has been shown to affect the market behavior significantly (Salganik, Dodds, and Watts, 2006; Lerman and Hogg, 2014; Abeliuk et al., 2015). Hence it is interesting to study how these recent results are affected by more realistic settings in which consumers can try multiple products. Also, experimental analysis using eye-tracking inspired cascade models, introduced first by Craswell et al. (2008), and showed fit experimental data better than separable models in presence of position bias. The intuition behind this model, is that users consider products in a top to bottom fashion, as they were presented in an ordered list, and they only look to the next product if the current one was not selected.

*Industrial and Operations Engineering & Computer Science and Engineering, University of Michigan, Ann Arbor (pvanhent@umich.edu).

†College of Engineering & Computer Science, Australian National University.

‡Melbourne Business School, The University of Melbourne.

This paper is an attempt at generalizing Multinomial Logit Models to account for a richer class of customer behavior. It endows the Multinomial Logit Model with a notion of continuation, which enables participants to sample multiple products before making a purchase. The paper studies how this generalization affects market efficiency and the role of social influence. The main contributions of the paper can be summarized as follows:

1. We show that a trial-offer market with continuation can be reduced to a traditional trial-offer market by adjusting the quality and appeal of the products and we quantify how the continuation model affects market efficiency;

2. We show that, under a natural continuation model, the quality-ranking policy, where the products are ranked by quality, is preserved by the reduction, but not the performance ranking, which optimizes the market performance at each step. We also show that social influence remains beneficial in this setting under the quality ranking;

3. Finally, we show experimental results that indicate that the popularity ranking, which ranks the product by popularity, benefits more from the generalization than the quality and performance ranking, unless the continuation is strongly dependent of the product just sampled. This improvement however is not enough to bridge the gap with the performance and quality rankings.

2 Trial-Offer Markets

This paper considers trial-offer markets in which participants can try a product before deciding whether to buy it. Such settings are common in online cultural markets (e.g., books, songs, and videos). In this paper, the trial-offer market is composed of n products and each product $i \in \{1, \dots, n\}$ is characterized by two values:

1. Its *appeal* A_i representing the inherent preference of trying product i;

2. Its *quality* q_i representing the probability of purchasing product i given that it was tried.

Each participant, when entering the market, is presented with a product list π: She then tries a product s in π and decides whether to purchase s with a certain probability. The product list is a permutation of $\{1, \dots, n\}$ and each position p in the list is characterized by its *visibility* $v_p > 0$ which is the inherent probability of trying a product in position p. Since the list π is a bijection from positions to products, its inverse is well-defined and is called a ranking. We denote rankings by σ in the following, π_i denotes the product in position i of the list π, and σ_i denotes the *position* of product i in the ranking σ. Therefore v_{σ_i} denotes the visibility of the position of product i.

The probability of trying product i given a list σ is

$$p_i(\sigma) = \frac{v_{\sigma_i} A_i}{\sum_{j=1}^{n} v_{\sigma_j} A_j}.$$

Given a ranking σ, the expected number of purchases is

$$\lambda(\sigma) = \sum_{i=1}^{n} p_i(\sigma)\, q_i. \tag{1}$$

The traditional static market optimization problem consists of finding a ranking σ^* maximizing $\lambda(\sigma)$, i.e.,

$$\sigma^* = \underset{\sigma \in S_n}{\text{arg-max}} \sum_{i=1}^{n} p_i(\sigma)\, q_i \tag{2}$$

where S_n represents the symmetry group over $\{1, \ldots, n\}$. Observe that consumer choice preferences for trying the products are essentially modeled as a discrete choice model based on a multinomial logit Luce (1965) in which product utilities are affected by their position.

Social Influence Following Krumme et al. (2012), this paper considers a dynamic market where the appeal of each product changes over time according to a social influence signal. Given a social signal $d = (d_1, \ldots, d_n)$, where d_i denotes the number of purchases of product i, the appeal of i becomes $A_i + d_i$ and hence the probability of trying i given a list σ becomes

$$p_i(\sigma, d) = \frac{v_{\sigma_i}(A_i + d_i)}{\sum_{j=1}^{n} v_{\sigma_j}(A_j + d_j)}.$$

Note that the probability of trying a product depends on its position in the list, its appeal, and its number of purchases ($d_{i,t}$) at time t. As the market evolves over time, the number of purchases could dominate the appeal, and the sampling probability of a product becomes its market share. Without social influence, a dynamic market reduces to solving the static optimization problem repeatedly. This set-up is the *independent condition*.

In the following, without loss of generality, we assume that the qualities and visibilities are non-increasing, i.e., $q_1 \geq q_2 \geq \cdots \geq q_n$ and $v_1 \geq v_2 \geq \cdots \geq v_n$. We also assume that the qualities and visibilities are known. In practical situations, the product qualities are obviously not known. But, as shown by Abeliuk et al. (2015), they can be recovered accurately and quickly, either before or during the market execution. For simplicity, we use $a_{i,t} = A_i + d_{i,t}$ to denote the appeal of product i at step t. When the step t is not relevant, we omit it and use a_i instead.

Ranking policies Following Abeliuk et al. (2015), this paper explores several ranking policies. The *performance ranking* maximizes the expected number of purchases at each iteration, exploiting all the available information globally, i.e., the appeal, the visibility, the purchases, and the quality of the products. More precisely, the performance ranking at step k produces a ranking σ_k^* defined as

$$\sigma_k^* = \underset{\sigma \in S_n}{\text{arg-max}} \sum_{i=1}^{n} p_i(\sigma, d_k) \cdot q_i$$

where $d_k = (d_{1,k}, \ldots, d_{n,k})$ is the social influence signal at step k. The performance ranking uses the probability $p_i(\sigma, d_k)$ of trying products i at iteration k given ranking σ, as well as the quality q_i of product i. The performance ranking can be computed in strongly polynomial time and the resulting policy is scalable to large markets Abeliuk et al. (2015). The *quality ranking* simply orders the products by quality, assigning the product of highest quality to the most visible position and so on. With the above assumptions, a quality ranking σ satisfies $\sigma_i = i \quad (1 \leq i \leq n)$. The *popularity ranking* was used by Salganik, Dodds, and Watts (2006) to show the unpredictability caused by social influence in cultural markets. At iteration k, the popularity ranking orders the products by the number of purchases $d_{i,k}$, but these purchases do not necessarily reflect the inherent quality of the products, since they depend on how many times the products were tried. We also follow Abeliuk et al. (2015) and use Q-RANK, D-RANK, and P-RANK to denote the policies using the quality, popularity, and performance rankings respectively. We also use R-RANK to denote the policy that simply presents a random order at each period.

3 Trial-Offer Markets With Continuation

The main goal of this paper is to study trial-offer markets with continuation, i.e., a setting where market participants can continue shopping even when they decline to purchase the product just sampled. We model such a trial-offer market by adding a continuation probability

$$c_i = f(\cdot)(1 - q_i) \tag{3}$$

to continue shopping after a participant has declined to purchase product i. In the above probability, the $(1 - q_i)$ term represents the fact that the participant has declined to purchase product i and the $f(\cdot)$ term represents a function that might depend on the product quality, the current position, or even on another overall measure (or a combination of all these factors). Figure 1 shows a graphic representation of a trial-offer market with continuation. It uses $\bar{c}_i = 1 - c_i$ to denote the probability that a participant leaves the market place after sampling product i.

The expected number of purchases in the static version of the trial-offer market with continuation for a ranking σ is denoted by $\overline{\lambda(\sigma)}$ and defined by

$$\overline{\lambda(\sigma)} = \sum_{i=1}^{n} p_i(\sigma)(q_i + c_i\overline{\lambda(\sigma)}) \tag{4}$$

Our primary objective is to maximize market efficiency, i.e., the expected purchases:

$$\sigma^* = \underset{\sigma \in S_n}{\text{arg-max}}\, \overline{\lambda(\sigma)}. \tag{5}$$

Note that the higher this objective is, the lower the probability that consumers try a product but then decide not to purchase it. Hence, if we interpret this last action as an inefficiency, maximising the expected efficiency of the market minimizes unproductive trials.

This paper proves a number of results when the continuation c_i depends polynomially on q_i, i.e.,

$$c_i = \rho q_i^r (1 - q_i) \tag{6}$$

where $\rho \leq 1$ controls the overall tendency to continuation and $r \geq 0$ represents the influence of q. This choice is justified intuitively by the fact that a market participant is more likely to continue sampling if the product she tried is of high quality, because it reflects on how good the other products potentially are. Figure 2 depicts various choices of ρ and r.

4 Reduction to the Trial-Offer Model

This section shows that the trial-offer market with continuation can be reduced to a trial-offer market. Indeed, rearranging the terms in Equation 4 leads to

$$\overline{\lambda(\sigma)} = \sum_{i=1}^{n} p_i(\sigma)(q_i + c_i\overline{\lambda(\sigma)})$$

$$\overline{\lambda(\sigma)} = \frac{\sum_{i=1}^{n} p_i(\sigma) q_i}{1 - \sum_{i=1}^{n} p_i(\sigma) c_i}$$

Defining

$$\overline{p_i(\sigma)} = \frac{p_i(\sigma)}{1 - \sum_{i=1}^{n} p_i(\sigma) c_i} \tag{7}$$

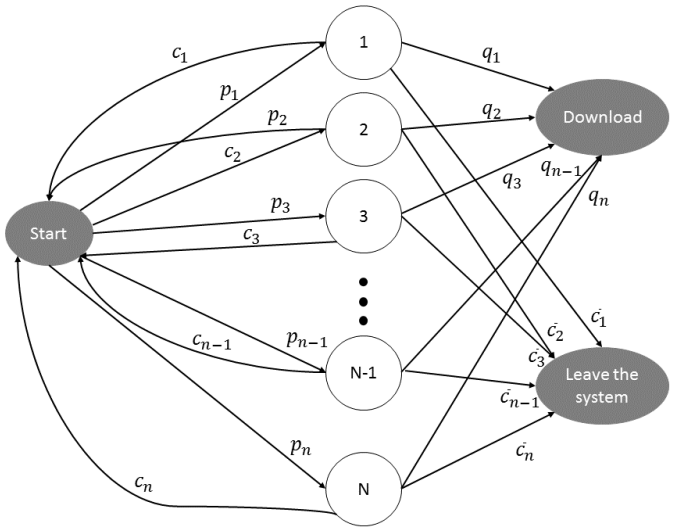

Figure 1: A Trial-Offer Market with Continuation.

we obtain

$$\overline{\lambda(\sigma)} = \sum_{i=1}^{n} \overline{p_i(\sigma)}\, q_i$$

By definition of $p_i(\sigma)$, we have

$$\overline{p_i(\sigma)} = \frac{v_{\sigma_i} a_i}{\sum_{i=1}^{n} v_{\sigma_i} a_i} \cdot \frac{1}{1 - \sum_{i=1}^{n} (c_i \cdot \frac{v_{\sigma_i} a_i}{\sum_{i=1}^{n} v_{\sigma_i} a_i})}$$
$$\overline{p_i(\sigma)} = \frac{v_{\sigma_i} a_i}{\sum_{i=1}^{n} (1 - c_i) v_{\sigma_i} a_i}$$

Now, by defining $\overline{a_i} = a_i(1 - c_i)$ and $\overline{q_i} = \frac{q_i}{(1-c_i)}$, we obtain

$$\overline{\lambda(\sigma)} = \sum_{i=1}^{n} \frac{v_{\sigma_i} \overline{a_i}\ \overline{q_i}}{\sum_{i=1}^{n} v_{\sigma_i} \overline{a_i}}$$

We have proven the following theorem:

Theorem 4.1. *A trial-offer market with continuation can be reduced to a trial-offer market by using the product qualities $\overline{q_i}$ and appeals $\overline{a_i}$ defined as follows:*

$$\overline{q_i} = \frac{q_i}{1 - c_i}$$
$$\overline{a_i} = a_i(1 - c_i).$$

In the following, $\overline{q_i}$ and $\overline{a_i}$ are called the *continuation qualities* and *continuation appeals*, and Figure 3 depicts the continuation quality for different values of ρ and r. Observe how the continuation model typically

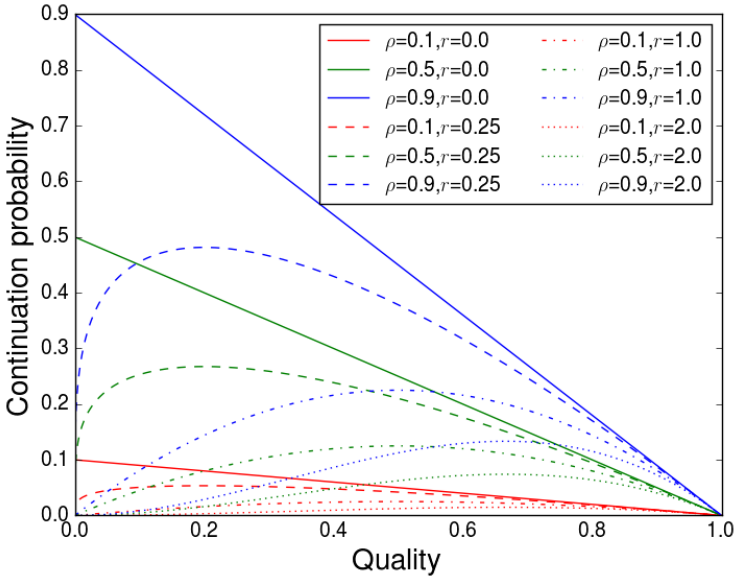

Figure 2: Examples of the continuation probabilities for different values of ρ and r; the r parameter defines where the peak is (the maximum is always attained at $q = \frac{r}{r+1}$), and ρ modulates how strong the continuation is.

boosts the quality of the products, sometimes substantially. To understand this reduction intuitively, we can rewrite Equation 7 as:

$$\overline{p_i(\sigma)} = p_i(\sigma) \cdot \sum_{j=1}^{\infty} (\sum_{i=1}^{n} p_i(\sigma) c_i)^j$$

The value $\overline{p_i(\sigma)}$ can thus be interpreted as the probability of sampling product i in any number of steps. The rewriting uses the fact that $\sum_{i=1}^{n} p_i(\sigma) c_i < 1$ to obtain an infinite sum and the term $(\sum_{i=1}^{n} p_i(\sigma) c_i)^j$ captures all the possible ways to sampling i in j steps.

5 Properties of the Market

Market Efficiency: The first result links the expected number of purchases of the market with and without continuation under the performance ranking.

Theorem 5.1. *Let π_c^* and π^* be optimal permutations for the trial-offer markets with and without continuation. Then,*

$$\lambda(\pi^*) \leq \overline{\lambda}(\pi_c^*) \leq \frac{\lambda(\pi^*)}{1 - \max_i c_i}.$$

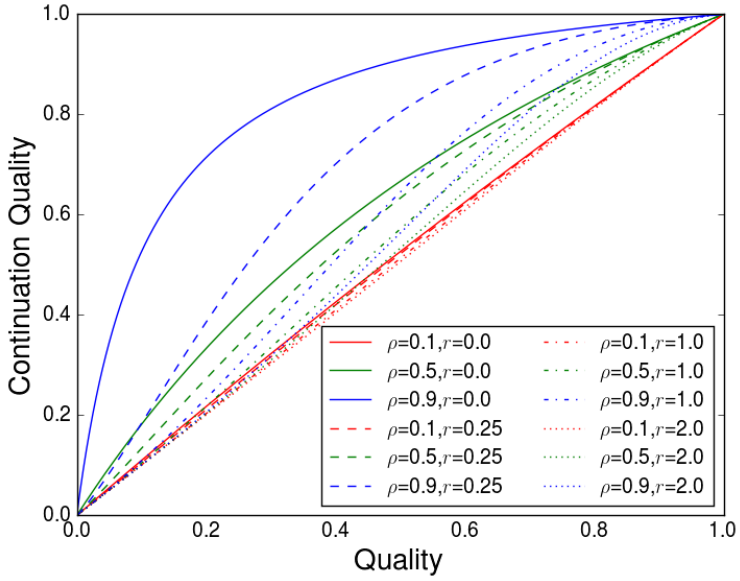

Figure 3: Continuation qualities for different values of the ρ and r parameters; the lager ρ is, the more *concave* the continuation quality becomes.

Proof. The lower bound can be derived as follows:

$$\lambda(\pi^*) = \frac{\sum_{i=1}^n v_i a_{\pi_i^*} q_{\pi_i^*}}{\sum_{i=1}^n v_i a_{\pi_i^*}}$$

$$\lambda(\pi^*) = \frac{\sum_{i=1}^n v_i a_{\pi_i^*} q_{\pi_i^*}}{\sum_{i=1}^n v_i a_{\pi_i^*}} \cdot \frac{1 - \sum_{i=1}^n p_i(\pi^*) c_{\pi_i^*}}{1 - \sum_{i=1}^n p_i(\pi^*) c_{\pi_i^*}}$$

$$\lambda(\pi^*) = \underbrace{\frac{\sum_{i=1}^n p_i(\pi^*) q_{\pi_i^*}}{1 - \sum_{i=1}^n p_i(\pi^*) c_{\pi_i^*}}}_{\overline{\lambda}(\pi^*)} \underbrace{(1 - \sum_{i=1}^n p_i(\pi^*) c_{\pi_i^*})}_{\leq 1}$$

$$\lambda(\pi^*) \leq \overline{\lambda}(\pi^*) \leq \overline{\lambda}(\pi_c^*)$$

where the last inequality holds because of optimality of π_c^* for $\overline{\lambda}(\cdot)$. The upper bound follows from

$$\overline{\lambda}(\pi_c^*) \leq \lambda(\pi_c^*) \cdot \frac{1}{1 - \sum_{i=1}^n p_i c_{\pi_i^*}} \leq \lambda(\pi^*) \cdot \frac{1}{1 - \max_i c_i}.$$

□

The following corollary considers the case where continuations depend polynomially on qualities (a proof is provided in the Appendix).

Corollary 5.1. *Assume that* $c_i = \rho q_i^r (1 - q_i)$. *It follows that*

$$\lambda(\pi^*) \leq \overline{\lambda}(\pi_c^*) \leq \lambda(\pi^*) \frac{1}{1 - \frac{\rho r^r}{(r+1)^{r+1}}} \tag{8}$$

When $\rho = 1$ and $r = 1$, $\lambda(\pi_c^*) \leq \frac{4}{3} \cdot \lambda(\pi^*)$ indicating a market that is at most 33% more efficient.

Prior work on trial-offer markets with the social influence signal considered here has shown that the quality ranking is asymptotically optimal Van Hentenryck et al. (2015): The market converges towards a monopoly for the product of highest quality. We now show that, when the continuations are polynomial in product qualities, the quality ranking is preserved by the reduction and hence the two markets, with and without continuation, converge to the same equilibrium in market shares.

Proposition 5.2. *Let $c_i = \rho q_i^r(1 - q_i)$ with $\rho \in (0, 1)$ and $r \geq 0$. Then $q_i \leq q_j \Leftrightarrow \overline{q_i} \leq \overline{q_j}$.*

Proof. It is sufficient to show that $\overline{q_i}$, when viewed as a function of q_i, is increasing in $(0, 1)$. Consider such function

$$h(x) = \frac{x}{1 - \rho x^r(1 - x)}$$

and its derivative

$$\frac{dh(x)}{dx} = \frac{\rho x^r \left[(r - 1) - rx\right] + 1}{((1 - \rho x^r(1 - x))^2}.$$

The denominator is greater than zero, so it remains to show that the numerator also is. The term ρx^r is increasing in x and the term $[(r - 1) - rx]$ is a line decreasing in x. The product is minimized when $x = 1$, in which case the product has a value of $-\rho$. Since $\rho \in (0, 1)$, the minimum value of the numerator is $1 - \rho \geq 0$, which concludes the proof. □

More importantly, it is also possible to show that, under the quality ranking, the probability that the next purchase is product i is the same in the markets with and without continuation. Hence, from a product standpoint, the markets behave very similarly.

Proposition 5.3. *The probability p_i that the next purchase (after any number of steps) is product i is*

$$p_i = \frac{v_i a_i q_i}{\sum_{j=1}^{n} v_j a_j q_j}.$$

A proof of Propositon 5.3 is given in the Appendix). In contrast, the same results do not hold for the performance ranking, which may change when a continuation is used, as shown by the following example.

Example 5.1. *Consider the following instance with 3 songs:*

- *Visibilities:* $v_1 = 0.8$, $v_2 = 0.5$ *and* $v_3 = 0.1$
- *Qualities:* $q_1 = 0.9$, $q_2 = 0.2$ *and* $q_3 = 0.6$
- *Appeals:* $a_1 = 0.9$, $a_2 = 0.1$ *and* $a_3 = 0.3$
- *Continuation parameters:* $\rho = 0.8$ *and* $r = 0.7$

In this case, the performance ranking for the market without continuation is $\sigma^* = [1, 2, 3]$; It is $\sigma_c^* = [1, 3, 2]$ for the continuation model.

Position Bias: This result generalizes the result shown in Van Hentenryck et al. (2015), to the continuation setting, and it means that we can always benefit from position bias. The formalization of this claim can be seen below (a proof is provided in the Appendix))

Theorem 5.2. *Position bias increases the expected number of purchases under the quality-ranking policy, i.e., for all visibilities v_i, appeals a_i, qualities q_i $(1 \leq i \leq n)$ and continuation probabilities c_i.This is, after we make the reduction to the Associated Multinomial Logit, we have:*

$$\frac{\sum_i v_i \overline{a_i}\ \overline{q_i}}{\sum_j v_j \overline{a_j}} \geq \frac{\sum_i \overline{a_i}\ \overline{q_i}}{\sum_j \overline{a_j}}.$$

Social Influence: The last result in this section shows that the social influence signals always benefit trial-offer markets with continuation. The result is independent of the structure of the continuation probabilities. The proof is a generalization of the result in Van Hentenryck et al. (2015).

Theorem 5.3. *The expected marginal rate of purchases is non-decreasing over time for the quality ranking under social influence in trial-offer markets with continuation.*

Proof. Let

$$\mathbb{E}[D_t] = \frac{\sum_i v_i \overline{a_i q_i}}{\sum_i v_i \overline{a_i}} = \overline{\lambda} \tag{9}$$

be the expected number of purchases at time t. The expected number of purchases at time $t+1$ conditional to time t is

$$\begin{aligned}
\mathbb{E}[D_{t+1}] &= \sum_j \left[\frac{v_j \overline{a_j q_j}}{\sum v_i \overline{a_i}} \cdot \frac{\sum_{i \neq j} v_i \overline{a_i q_i} + v_j(\overline{a_j} + 1 - c_j)\overline{q_j}}{\sum_{i \neq j} v_i \overline{a_i} + v_j(\overline{a_j} + 1 - c_j)} \right] \\
&\quad + \left[1 - \frac{\sum_i v_i \overline{a_i q_i}}{\sum_i v_i \overline{a_i}} \right] \cdot \frac{\sum_i v_i \overline{a_i q_i}}{\sum_i v_i \overline{a_i}} \\
&= \sum_j \left[\frac{v_j \overline{a_j q_j}}{\sum v_i \overline{a_i}} \cdot \frac{\sum_i v_i \overline{a_i q_i} + v_j(1 - c_j)\overline{q_j}}{\sum_i v_i \overline{a_i} + v_j(1 - c_j)} \right] \\
&\quad + \left[1 - \frac{\sum_j v_j \overline{a_j q_j}}{\sum_i v_i \overline{a_i}} \right] \cdot \overline{\lambda}
\end{aligned}$$

We need to prove that

$$\mathbb{E}[D_{t+1}] \geq \mathbb{E}[D_t], \tag{10}$$

which is equivalent to show, using Equation 9, that

$$\sum_j \left[\frac{v_j \overline{a_j q_j}}{\sum v_i \overline{a_i}} \cdot \frac{\sum_i v_i \overline{a_i q_i} + v_j(1 - c_j)\overline{q_j}}{\sum_i v_i \overline{a_i} + v_j(1 - c_j)} \right] + [1 - \overline{\lambda}] \cdot \overline{\lambda} \geq \overline{\lambda}$$

Rearranging the terms, the proof obligation becomes

$$\frac{1}{\sum_i v_i \overline{a_i}} \sum_j \left[\frac{v_j^2 \overline{a_j q_j}(1 - c_j)}{\sum_i v_i \overline{a_i} + v_j(1 - c_j)} \left(\overline{q_j} - \overline{\lambda} \right) \right] \geq 0$$

or, equivalently,

$$\sum_j \left[\frac{v_j^2 \overline{a_j q_j}(1-c_j)}{\sum_i v_i \overline{a_i} + v_j(1-c_j)} \left(\overline{q_j} - \overline{\lambda}\right) \right] \geq 0. \tag{11}$$

Let $k = \max\{i \in N | (\overline{q_i} - \lambda) \geq 0\}$, i.e., the largest $k \in N$ such that $q_k \geq \lambda$. By separating the sum into positive and negative terms, we obtain

$$\sum_j \left[\frac{v_j^2 \overline{a_j q_j}(1-c_j)\left(\overline{q_j} - \overline{\lambda}\right)}{\sum_i v_i \overline{a_i} + v_j(1-c_j)} \right] = S^+ + S^- \quad \text{where}$$

$$S^+ = \sum_{j=1}^{k} \left[\frac{v_j \overline{q_j}(1-c_j)}{\sum_i v_i \overline{a_i} + v_j(1-c_j)} \overline{a_j} v_j (\overline{q_j} - \overline{\lambda}) \right],$$

$$S^- = \sum_{j=k+1}^{n} \left[\frac{v_j \overline{q_j}(1-c_j)}{\sum_i v_i \overline{a_i} + v_j(1-c_j)} \overline{a_j} v_j (q_j - \overline{\lambda}) \right]$$

By definition of k, all the terms in S^+ are positive and the terms in S^- are negative. Now, by definition of k and $\overline{q_i} = \frac{q_i}{(1-c_i)}$, we have

$$\forall i \leq k: \ (1-c_i) \leq \frac{q_i}{\overline{\lambda}},$$
$$\forall i > k: \ (1-c_i) \geq \frac{q_i}{\overline{\lambda}}. \tag{12}$$

We now compute a lower bound for S^+ and S^-. For S^+, using Equation 12 for $j \leq k$, we have

$$\frac{v_j \overline{q_j}(1-c_j)}{\sum_i v_i \overline{a_i} + v_j(1-c_j)} \leq \frac{v_j q_j}{\sum_i v_i \overline{a_i} + v_j \frac{q_j}{\overline{\lambda}}}$$
$$\leq \overline{\lambda} \cdot \frac{v_j q_j}{\overline{\lambda} \sum_i v_i \overline{a_i} + v_j q_j} \leq \overline{\lambda} \cdot \frac{v_k q_k}{\overline{\lambda} \sum_i v_i \overline{a_i} + v_k q_k} \tag{13}$$

The last inequality follows by $v_i \geq v_k$ and $q_i \geq q_k$ (using Theorem 5.2) and the following property: For all $c > 0$ and $x \geq y \geq 0$,

$$\frac{x}{c+x} \geq \frac{y}{c+y} \Leftrightarrow (c+y)x \geq (c+x)y \Leftrightarrow cx \geq cy \Leftrightarrow x \geq y.$$

For S^-, consider the following expression for $j > k$:

$$\overline{\lambda}(\sum_i v_i \overline{a_i}) \underbrace{[v_j q_j - v_k q_k]}_{\geq 0} + v_k q_k \underbrace{[v_j q_j - \overline{\lambda} v_j (1-c_j)]}_{\geq 0}$$

Here the first term is greater or equal than zero because $v_j \geq v_k$ and $q_j \geq q_k$ using Theorem 5.2 again. The second term is also greater than zero because it can be lower-bounded (using Equation 12) by:

$$v_j q_j - \overline{\lambda} \frac{v_j q_j}{\overline{\lambda}} = 0.$$

Hence,

$$\overline{\lambda}(\sum_i v_i\overline{a_i})[v_jq_j - v_kq_k] + v_kq_k[v_jq_j - \overline{\lambda}v_j(1-c_j)] \geq 0$$

$$v_jq_j[\overline{\lambda}(\sum_i v_i\overline{a_i}) + v_kq_k] \geq \overline{\lambda}v_kq_k[\sum_i v_i\overline{a_i} + v_j(1-c_j)]$$

$$\Leftrightarrow \frac{v_jq_j}{\sum_i v_i\overline{a_i} + v_j(1-c_j)} \geq \frac{\overline{\lambda}v_kq_k}{\overline{\lambda}\sum_i v_i\overline{a_i} + v_kq_k} \tag{14}$$

Putting together Equations 13 and 14 gives us a lower bound to $S^+ + S^-$:

$$S^+ + S^- = \frac{\overline{\lambda}v_kq_k}{\overline{\lambda}\sum_i v_i\overline{a_i} + v_kq_k} \cdot \sum_{i=1}^{n} v_i\overline{a_i}(\overline{q_i} - \overline{\lambda}) \tag{15}$$

Now, by definition of $\overline{\lambda}$,

$$\overline{\lambda} = \frac{\sum_{i=1}^n v_i\overline{a_iq_i}}{\sum_{i=1}^n v_i\overline{a_i}} \Leftrightarrow \sum_{i=1}^{n} v_i\overline{a_i}(\overline{q_i} - \overline{\lambda}) = 0.$$

which implies that

$$\frac{\overline{\lambda}v_kq_k}{\overline{\lambda}\sum_i v_i\overline{a_i} + v_kq_k} \cdot \sum_{i=1}^{n} v_i\overline{a_i}(\overline{q_i} - \overline{\lambda}) = 0$$

concluding the proof. □

Relationship with the Cascade Model: Observe that the quality ranking over the continuation quality orders the products in decreasing order of $\frac{q_i}{1-c_i}$ value, which is exactly the adjusted ecpm from Aggarwal et al. (2008); Kempe and Mahdian (2008) with all the revenues set to 1. Obviously, the quality ranking (when the continuation probabilities preserve the quality rank in the continuation model), and hence the adjusted ecpm ranking, are not the best rankings to show to an incoming participant (the performance ranking is), but our results show that they have nice asymptotic properties.

6 Experimental Results

This section report computational results to highlight the theoretical analysis. The computational results use settings that model the MUSICLAB experiments discussed in Salganik, Dodds, and Watts (2006); Krumme et al. (2012); Abeliuk et al. (2015). As mentioned in the introduction, MUSICLAB is a trial-offer market where participants can try a song and then decide to download it. The experiments use an agent-based simulation to emulate MUSICLAB. Each simulation consists of N steps and, at each iteration t,

1. we simulate selecting a song i according to the probabilities $p_i(\sigma, d)$, where σ is the ranking proposed by the policy under evaluation and d is the social influence signal.

2. with probability q_i, the sampled song is downloaded, in which case the simulator increases the social influence signal for song i, i.e., $d_{i,t+1} = d_{i,t} + 1$. Otherwise, $d_{i,t+1} = d_{i,t}$, and if the continuation model is used, the simulation goes back to Step 1 with probability c_i and advances to the next step otherwise.

Every T iterations, a new list σ is computed using one of the ranking policies. The experimental setting, which aims at being close to the MUSICLAB experiments, considers 50 songs and simulations with 20,000 steps. The songs are displayed in a single column. The analysis in Krumme et al. (2012) indicated that participants are more likely to try songs higher in the list. More precisely, the visibility decreases with the list position, except for a slight increase at the bottom positions. This paper uses the first setting for qualities, appeals and visibilities from Abeliuk et al. (2015), where the quality and the appeal are chosen independently according to a Gaussian distribution normalized to fit between 0 and 1. In addition, the experiments consider 12 different continuation probabilities, varying ρ and the power r as shown in Figure 3. The results were obtained by averaging $W = 100$ simulations.

Parameters	P-RANK	Q-RANK	D-RANK	R-RANK
$\rho = 0.1, r = 0$	5.3%	4.9%	5.8%	7.5%
$\rho = 0.1, r = 0.25$	4.1%	4.3%	4.9%	5.2%
$\rho = 0.1, r = 1$	2.2%	2.4%	2.6%	2.5%
$\rho = 0.1, r = 2$	1.4%	1.4%	0.2%	1%
$\rho = 0.5, r = 0$	30.6%	31%	38.3%	51.8%
$\rho = 0.5, r = 0.25$	24.2%	24.6%	28.4%	33.6%
$\rho = 0.5, r = 1$	13.4%	13.2%	14.8%	12.2%
$\rho = 0.5, r = 2$	7.2%	7.3%	6.2%	4.6%
$\rho = 0.9, r = 0$	67.3%	67.7%	93.9%	143.3%
$\rho = 0.9, r = 0.25$	51.6%	52.1%	65.2%	79.9%
$\rho = 0.9, r = 1$	26.6%	26.8%	28.5%	24.2%
$\rho = 0.9, r = 2$	13.6%	13.7%	12.2%	8.3%

Table 1: Improvement in Market Efficiency (in percentage) for the Continuation Model.

Parameters	P-RANK	Q-RANK	D-RANK	R-RANK
$\rho = 0.5, r = 0.25$	13776.1	13804.1	12000.1	9393.8
$\rho = 0.5, r = 1$	12579.0	12565.5	10643.7	7885.0
$\rho = 0.9, r = 0.25$	16784.7	16840.9	15435.7	12680.8
$\rho = 0.9, r = 1$	14041.4	14059.1	11926.5	8741.7

Table 2: Market Efficiency in the Continuation Model.

Table 1 presents results on market efficiency (i.e., the number of downloads) for the trial and offer market with continuation. The most interesting message from these results, is the observation that the popularity and random rankings improve more than the performance and quality rankings, unless the quality has less impact ($r = 2$, because the continuation is decreasing in r) in the continuation. This can be explained by the fact that the continuation provides a way to correct a potentially weak ranking. However, as indicated in Table 2, this correction is not enough to bridge the gap with the performance and quality rankings.

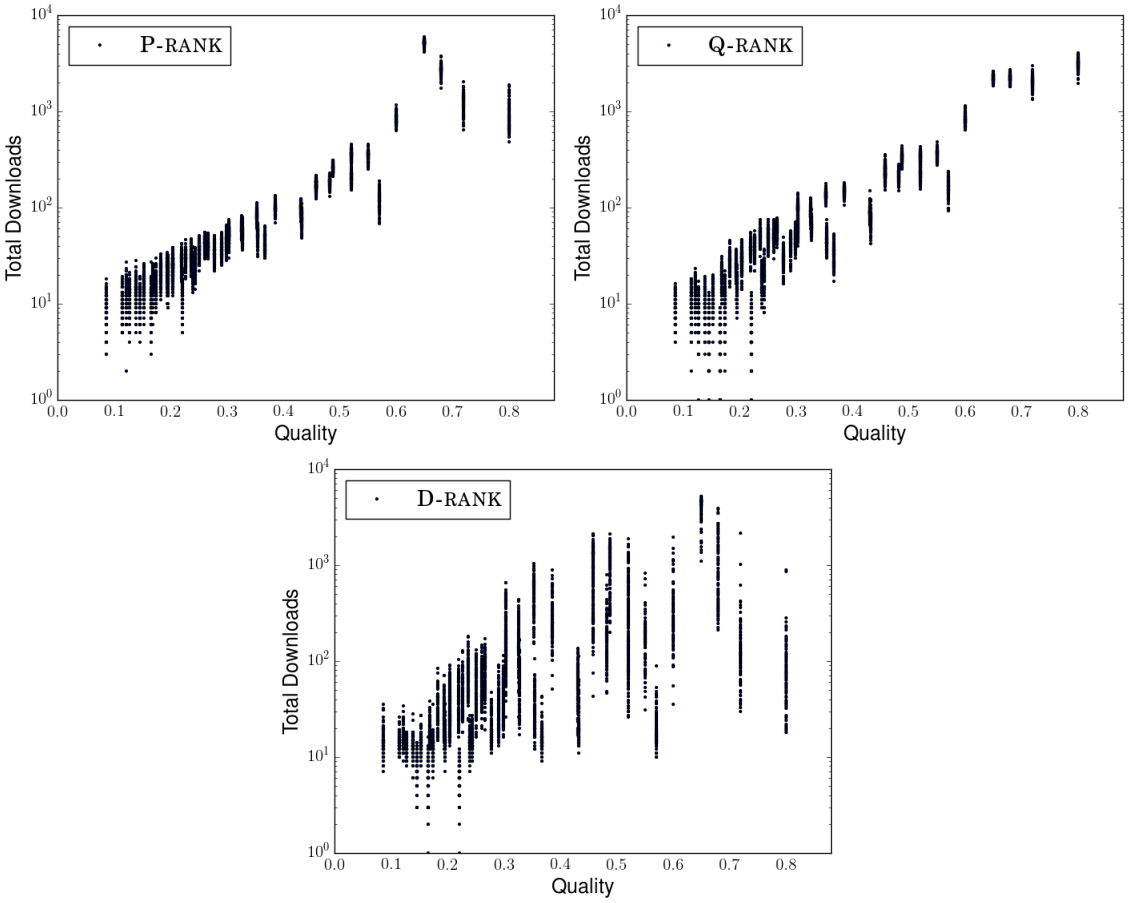

Figure 4: The Distribution of Downloads Versus Song Qualities for $\rho = 0.9$, $r = 1$. The songs on the x-axis are ranked by increasing quality from left to right. Each dot is the number of download of a product in one of the 100 experiments.

Figure 4 depicts experimental results on the predictability of the market under the continuation model under various ranking policies. The figure plots the number of downloads of each song for 100 experiments. In the plots, the songs are ranked by increasing quality from left to right on the x-axis. Each dot in the plot shows the number of downloads of a song in one of the 100 experiments. The results are essentially unchanged when moving from the traditional to a continuation multinomial logit model. The popularity ranking still exhibits significantly more unpredictability than the performance and quality rankings and the continuations are not able to compensate for the inherent unpredictability.

7 Conclusion and Future Work

Motivated by applications in online markets, this paper generalizes the ubiquitous multinomial logit model to a setting that allows market participants to sample multiple products before deciding whether to purchase. The paper showed that trial-offer markets with continuation can be reduced to the original trial-offer model, transferring many fundamental properties of ranking policies to a more general setting. In particular, the quality ranking still benefits from position bias and social influence. Moreover, under a general class of continuations, the quality ranking is also preserved and the market reaches the same asymptotic equilibrium.

Experimental results shows that the continuation model compensates for some of the weaknesses of the popularity ranking by boosting its market performance more than the quality and performance ranking, unless the continuation probability depends too strongly on quality. Our current research aims at generalizing these results further to hierarchical trial-offer markets.

References

Abeliuk, A.; Berbeglia, G.; Cebrian, M.; and Van Hentenryck, P. 2015. The benefits of social influence in optimized cultural markets. *PLoS ONE* 10(4):e0121934.

Aggarwal, G.; Feldman, J.; Muthukrishnan, S.; and Pal, M. 2008. Sponsored search auctions with markovian users. *CoRR* abs/0805.0766.

Buscher, G.; Cutrell, E.; and Morris, M. R. 2009. What do you see when you're surfing?: using eye tracking to predict salient regions of web pages. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 21–30. ACM.

Craswell, N.; Zoeter, O.; Taylor, M.; and Ramsey, B. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, WSDM '08, 87–94. New York, NY, USA: ACM.

Daly, A., and Zachary, S. 1978. Improved multiple choice models. *Determinants of travel choice* 335:357.

Engstrom, P., and Forsell, E. 2014. Demand effects of consumers' stated and revealed preferences. *Available at SSRN 2253859*.

Joachims, T.; Granka, L.; Pan, B.; Hembrooke, H.; and Gay, G. 2005. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, 154–161. ACM.

Kempe, D., and Mahdian, M. 2008. A cascade model for externalities in sponsored search. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, WINE '08, 585–596. Berlin, Heidelberg: Springer-Verlag.

Krumme, C.; Cebrian, M.; Pickard, G.; and Pentland, S. 2012. Quantifying social influence in an online cultural market. *PLOS ONE* 7(5):e33785.

Lerman, K., and Hogg, T. 2014. Leveraging position bias to improve peer recommendation. *PLOS ONE* 9(6).

Luce, R. D. 1965. *Individual Choice Behavior*. John Wiley and Sons.

Maillé, P.; Markakis, E.; Naldi, M.; Stamoulis, G. D.; and Tuffin, B. 2012. Sponsored search auctions: an overview of research with emphasis on game theoretic aspects. *Electronic Commerce Research* 12(3):265–300.

Rusmevichientong, P.; Shen, Z.-J. M.; and Shmoys, D. B. 2010. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research* 58(6):1666–1680.

Rusmevichientong, P.; Shmoys, D.; and Topaloglu, H. 2010. Assortment optimization with mixtures of logits. Technical report, School of IEOR, Cornell University.

Salganik, M. J.; Dodds, P. S.; and Watts, D. J. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762):854–856.

Talluri, K., and Van Ryzin, G. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1):15–33.

Van Hentenryck, P.; Abeliuk, A.; Berbeglia, F.; and Berbeglia, G. 2015. On the optimality and predictability of cultural markets with social influence. *CoRR* abs/1505.02469.

Viglia, G.; Furlan, R.; and Ladrón-de Guevara, A. 2014. Please, talk about it! when hotel popularity boosts preferences. *International Journal of Hospitality Management* 42:155–164.

Appendix

In this section we provide the proofs missing from the main text.

Proof of Corollary 5.1. The proof follows from examining the maximum value for the c_i.

$$\max_i c_i = \max_i \rho q_i^r (1 - q_i) \leq \rho \max_{x \in [0,1]} x^r (1 - x) = \frac{\rho r^r}{(r+1)^{r+1}}$$

where the last equality holds because the maximum value of $x^r(1-x)$ is reached when $x = \frac{r}{r+1}$. □

Proof of Corollary 5.1. The probability that product i is purchased in the first step is

$$p_i^{1st} = \frac{v_i a_i}{\sum_{j=1}^{n} v_j \overline{a_j}} q_i,$$

More generally, the probability that product i is purchased in step m while no product was purchased in earlier steps is:

$$p_i^{mth} = \left(\frac{\sum_{j=1}^{n} v_j \overline{a_j} (1 - \overline{q_j})}{\sum_{j=1}^{n} v_j \overline{a_j}} \right)^{m-1} \frac{v_i a_i}{\sum_{j=1}^{n} v_j \overline{a_j}} q_i.$$

Defining $\beta = (\sum_{j=1}^{n} v_j a_j q_j)/(\sum_{j=1}^{n} v_j \overline{a_j})$, Equation 7 becomes

$$p_i^{mth} = \Big(1 - \beta\Big)^{m-1} \frac{v_i a_i}{\sum_{j=1}^{n} v_j \overline{a_j}} q_i.$$

Hence the probability that the next purchased product is product i is given by

$$p_i = \sum_{m=0}^{\infty} \Big(1 - \beta\Big)^{m} \frac{v_i a_i}{\sum_{j=1}^{n} v_j \overline{a_j}} q_i.$$

Given the fact that $\beta < 1$ we have:

$$\sum_{m=0}^{\infty} \Big(1 - \beta\Big)^{m} = \frac{1}{\beta},$$

the probability that the next purchase is product i is given by

$$p_i = \frac{\sum_{j=1}^{n} v_j \overline{a_j}}{\sum_{j=1}^{n} v_j a_j q_j} \cdot \frac{v_i a_i}{\sum_{j=1}^{n} v_j \overline{a_j}} q_i.$$

$$p_i = \frac{v_i a_i q_i}{\sum_{j=1}^{n} v_j a_j q_j}$$

□

Proof of Theorem 5.2. Let $\overline{\lambda} = \frac{\sum_i v_i \overline{a_i q_i}}{\sum_j v_j \overline{a_j}}$ be the expected number of purchases for the quality ranking. We have

$$\sum_i v_i \overline{a_i} \left(\overline{q_i} - \overline{\lambda} \right) = 0.$$

Consider the index k such that $\left(\overline{q_k} - \overline{\lambda} \right) \geq 0$ and $(\overline{q_{k+1}} - \lambda) < 0$. Since $v_1 \geq \ldots \geq v_n$, we have

$$\sum_{i=1}^{k} v_k \overline{a_i} \left(\overline{q_i} - \overline{\lambda} \right) + \sum_{i=k+1}^{n} v_k \overline{a_i} \left(\overline{q_i} - \overline{\lambda} \right) \leq \sum_i v_i \overline{a_i} \left(\overline{q_i} - \overline{\lambda} \right) = 0$$

and, given the fact that $v_k \geq 0$,

$$\sum_{i=1}^{n} \overline{a_i} \left(\overline{q_i} - \overline{\lambda} \right) \leq 0.$$

We had the desired result: $\lambda \geq \frac{\sum_{i=1}^{n} \overline{a_i q_i}}{\sum_{i=1}^{n} \overline{a_i}}$. □